Steiner Point Removal — Distant Terminals Don't (Really) Bother

Yun Kuen Cheung
Max Planck Institute for Informatics
Saarland Informatics Campus

arXiv:1703.08790v1 [cs.DS] 26 Mar 2017

Abstract

Given a weighted graph $G = (V, E, w)$ with a set of k terminals $T \subset V$, the Steiner Point Removal problem seeks for a minor of the graph with vertex set T, such that the distance between every pair of terminals is preserved within a small multiplicative distortion. Kamma, Krauthgamer and Nguyen (SODA 2014, SICOMP 2015) used a ball-growing algorithm to show that the distortion is at most $\mathcal{O}(\log^5 k)$ for general graphs.

In this paper, we improve the distortion bound to $\mathcal{O}(\log^2 k)$. The improvement is achieved based on a known algorithm that constructs terminal-distance exact-preservation minor with $\mathcal{O}(k^4)$ (which is independent of $|V|$) vertices, and also two tail bounds on the sum of independent exponential random variables, which allow us to show that it is unlikely for a non-terminal being contracted to a distant terminal.

Keywords. Steiner Point Removal, Graph Sparsification, Vertex Sparsification, Exponential Random Variables

1 Introduction

Graph Compression/Sparsification generally describes a transformation of a *large* graph into a *smaller* graph that preserves, either exactly or approximately, certain features (e.g., distance, cut, flow) of the large graph. Its algorithmic value is clear, since the compressed graph can be computed in a preprocessing step of an algorithm, so as to reduce subsequent running time and memory requirement. Some notable examples are graph spanners, low-stretch spanning tree, distance oracles and cut/flow/spectral sparsifiers.

In this paper, we study a vertex sparsification problem called the *Steiner Point Removal (SPR) problem*. Given a weighted graph $G = (V, E, w)$ with a set of k terminals $T \subset V$, the SPR problem seeks for a new graph G' which is a minor of G, such that the vertex set of G' is T (i.e., no non-terminal in the minor) and the distance between every pair of terminals is preserved within a multiplicative distortion κ:

$$\forall t_i, t_j \in T, \quad \mathsf{dist}_G(t_i, t_j) \ \leq \ \mathsf{dist}_{G'}(t_i, t_j) \ \leq \ \kappa \cdot \mathsf{dist}_G(t_i, t_j).$$

The target is to minimize κ. The requirement that G' is a minor of G is crucial since minor operations preserve certain structural similarities, e.g., planarity, of the input graph G.

SPR problem was first initiated by Gupta [Gup01], where the input graphs are trees. He showed that for tree graphs, the distortion is at most 8. A matching lower bound was shown by Chan et al. [CXKR06]. Recently, Kamma, Krauthgamer and Nguyen[1] [KKN15] used a randomized ball-growing algorithm to show the interesting result that the distortion is at most $\mathcal{O}(\log^5 k)$ for general graphs. It remains a big gap between their upper bound and the best known lower bound (which is 8); narrowing this gap is the motivation of the current work.

[1] Since we will refer to this paper for multiple times, for brevity we will refer it as KKN.

In this paper, we improve the distortion bound to $\mathcal{O}(\log^2 k)$. Our algorithm starts with a preprocessing step, which uses an algorithm of Krauthgamer, Nguyen and Zondiner [KNZ14] to obtain a minor of the input graph such that all terminal distances are preserved exactly, while the minor contains only $\mathcal{O}(k^4)$ vertices. Then we use the randomized ball-growing algorithm of KKN (with a few adjustments on parameters) on the preprocessed graph. By working with the preprocessed graph that contains only $\mathsf{poly}(k)$ vertices, we can define a class of $\mathsf{poly}(k)$ "bad events" such that avoidances of all these bad events will lead to a minor with $\mathcal{O}(\log^2 k)$ distortion. Finally, we show that the algorithm of KKN avoids all bad events with high probability.

In KKN's analysis, they first proved the result for graphs with bounded diameters; for graphs with any diameters, they needed to provide a separate argument to reduce to low-diameter scenarios. In contrast, our analysis will work directly — we do not need any assumption on the graph diameters. Also, our analysis bypasses KKN's need for analyzing *long active subpaths*. Hence, our analysis is more compact than KKN's analysis.

We note that an important component of our analysis is two tail bounds on the sum of independent exponential random variables. This allows us to show that it is unlikely for a non-terminal being contracted to a *distant* terminal.

Further Related Work. Basu and Gupta [BG08] showed that for outer-planar graphs (planar graphs with all terminals on the outermost face), SPR problem can be solved with distortion $\mathcal{O}(1)$. When randomization is allowed, Englert et al. [EGK$^+$14] showed that for graphs that exclude a fixed minor, one can construct a randomized minor for SPR problem with $\mathcal{O}(1)$ *expected* distortion. It remains open on whether similar guarantees can be obtained in the deterministic setting.

A natural generalization of SPR problem is to allow the minor keeping a small number of non-terminals. This generalization was initiated by Krauthgamer, Nguyen and Zondiner [KNZ14], in which they focused on preserving terminal distances *exactly* for general graphs, trees, planar graphs and graphs with bounded treewidth. Cheung, Gramoz and Henzinger [CGH16] extended to the setting of preserving terminal distances *approximately*. They proved super-linear (in k) lower bounds on the number of non-terminals required in the minor for achieving distortion less than 8. They also extended the technique for proving the lower bounds to establish the following result related to SPR problem: if the tight bound of SPR problem is super-constant (in k), then allowing the freedom of keeping $\mathcal{O}(k)$ non-terminals in the minor would *not* improve the distortion bound to a constant.

In many other settings, problems that concern preserving features which involve only a small number of terminals are becoming more popular. Terminal distance oracles were studied by Roditty, Thorup and Zwick [RTZ05]. They showed that there exists oracle using only $\mathcal{O}(\tau \cdot |V| \cdot k^{1/\tau})$ space which answers the distance between a terminal and any other vertex with stretch at most $(2\tau - 1)$. Elkin, Filtser and Neiman [EFN15a] improved the space requirement to $\mathcal{O}(\tau \cdot k^{1/\tau} + |V|)$, with a slightly worse stretch of at most $(4\tau - 1)$. Recently, Elkin, Filtser and Neiman [EFN15b] considered the problem of embedding a finite metric (e.g., graph metric) into normed metric while preserving all distances from any terminal to all other vertices approximately. They showed that the distortion depends only on k, but not on $|V|$. For cut/flow sparsifiers, we refer readers to [LM10, Moi09, CLLM10, MM16, Chu12, KR13, AGK14] for more details.

2 Main Theorem and Preliminary

For any graph G, G' is a minor of G if G' can be attained from G by a sequence of edge contractions, edge deletions and vertex deletions.

Given a weighted connected graph $G = (V, E, w)$ with a set of k terminals $T = \{t_1, t_2, \cdots, t_k\} \subset V$, the Steiner Point Removal (SPR) problem seeks to construct a minor $G' = (T, E', w')$ of the graph G, such that for any $t_i, t_j \in T$, $\mathsf{dist}_G(t_i, t_j) \leq \mathsf{dist}_{G'}(t_i, t_j) \leq \kappa \cdot \mathsf{dist}_G(t_i, t_j)$ with the minimum possible distortion κ. Our main result is that for general graphs, the distortion is at most $\mathcal{O}(\log^2 k)$:

Theorem 1. *For any weighted graph with k terminals, the Steiner Point Removal problem can be solved with distortion at most $\mathcal{O}(\log^2 k)$. More precisely, when k is sufficiently large, the distortion is at most $(2 \times 10^8) \cdot \log^2 k$.*

We note that our analysis is rather loose on the constant factor, so there should be plenty of room for reducing the constant 2×10^8. But we have not yet done so for a cleaner analysis.

Notations. By using a consistent tie-breaking rule (e.g., edge weight perturbation), we can assume that there is a unique shortest path between any two vertices in G. We let $\mathsf{SP}(v_i, v_j)$ denote the unique *directed* shortest path in G from vertex v_i to vertex v_j.

For any $V' \subset V$, let $G[V']$ be the subgraph of G induced by V'. If $t \in V'$, let $B_{G[V']}(t, R) := \{v \in V' \mid \mathsf{dist}_{G[V']}(t, v) \leq R\}$.

Terminal-centered Minors. As KKN showed, SPR problem is equivalent to finding the *terminal-centered minor* of V that minimizes the quantity $\max_{t_i, t_j \in T} \frac{\mathsf{dist}_{G'}(t_i, t_j)}{\mathsf{dist}_G(t_i, t_j)}$. A terminal-centered minor of $G = (V, E, w)$ is specified by a *terminal-centered partition* of V into exactly k sets, which satisfies (a) each set contains exactly one of the terminals, and thus the set containing t_j can be naturally denoted by $V_j \equiv V_{t_j}$; (b) for each $j \in [k]$, the induced graph $G[V_j]$ is connected. Then the terminal centered minor $G' = (T, E', w')$ is formed by contracting the vertices in each V_j into a single vertex, which is identified with t_j. For any $t_i, t_j \in T$, $(t_i, t_j) \in E'$ if and only if there exists an edge $(u, v) \in E$ such that one of its endpoints belongs to V_i and the other endpoint belongs to V_j. For each $(t_i, t_j) \in E'$, its weight $w'((t_i, t_j)) := \mathsf{dist}_G(t_i, t_j)$.

Exponential Random Variables. An exponential random variable (ERV) with mean μ is the probability distribution with density function $f(x) = \mu^{-1} \cdot e^{-x/\mu}$ on $x \geq 0$, and $f(x) = 0$ on $x < 0$. We denote the distribution by $\mathsf{Exp}(\mu)$. ERV enjoys two important properties: (a) closeness under scaling: the random variable $\mathsf{Exp}(\mu)$ follows the same distribution as the random variable $\mu \cdot \mathsf{Exp}(1)$, and (b) memoryless property: for any $a, b \geq 0$, $\mathbb{P}[\, \mathsf{Exp}(\mu) > a + b \mid \mathsf{Exp}(\mu) > a \,] = \mathbb{P}[\, \mathsf{Exp}(\mu) > b \,]$. We will also use the following inequality: for any $\kappa \geq 0$, $\mathbb{P}[\, \mathsf{Exp}(\mu) \leq \kappa\mu \,] \leq \kappa$.

3 Algorithm and Technical Overview

Algorithm 1: The Ball-Growing Algorithm that produces a terminal-centered partition.

INPUT: A weighted graph $G = (V, E, w)$, and a set of terminals $T = \{t_1, \cdots, t_k\} \subset V$.
OUTPUT: A terminal-centered partition $\{V_1, \cdots, V_k\}$ of V.

Set $r \leftarrow 1 + \delta / \log k$, where $\delta = 1/2$.
Set $D \leftarrow \frac{\delta}{100 \log k} \cdot \min_{v \in V \setminus T} D_v$.
For each $j \in [k]$, set $V_j \leftarrow \{t_j\}$, and set $R_j \leftarrow 0$. *// $R_j \equiv R_{t_j}$.*
Set $V_\perp \leftarrow V \setminus \left(\cup_{j=1}^k V_j \right)$. *// $V_\perp$ keeps track of the set of unassigned vertices.*
Set $\ell \leftarrow 0$.
while $\left(\cup_{j=1}^k V_j \right) \neq V$ **do**
    //We call this round by Round-$(D \cdot r^\ell)$;
    forall $j \in [k]$ **do**
        Choose independently at random $q_j^\ell \sim \mathsf{Exp}(D \cdot r^\ell)$.
        Set $R_j \leftarrow R_j + q_j^\ell$.
        Set $V_j \leftarrow V_j \cup B_{G[V_\perp \cup V_j]}(t_j, R_j)$.
        Set $V_\perp \leftarrow V \setminus \left(\cup_{j=1}^k V_j \right)$.
    Set $\ell \leftarrow \ell + 1$.
return $\{V_1, \cdots, V_k\}$.

3.1 Algorithm

Our first step is to use an algorithm of Krauthgamer et al. [KNZ14, Theorem 2.1] to obtain a minor of G, denoted by G_*, such that all terminal distances are preserved exactly, and G_* contains only $\mathcal{O}(k^4)$ non-terminals — note that this number does not depend on $|V|$ and is in $\mathsf{poly}(k)$. In the rest of this paper, we identify G_* as G.

Next, we use the randomized ball-growing algorithm of KKN (with a few adjustments on parameters, see Algorithm 1) on the graph G to construct a terminal-centered partition. In the course of the algorithm, a vertex v is *assigned* to terminal j if $v \in V_j$; we say a vertex v is *assigned* if it is assigned to some terminal, and we say v is *unassigned* otherwise. As already noted in Algorithm 1, we always refer to a round by Round-μ, where μ is the mean of the ERV generated in that round. We abuse notation a bit: for any positive real number z, by Round-z we refers to the first Round-μ in the algorithm such that $\mu \geq z$; note that $\mu \leq \frac{3}{2}z$.

3.2 Families of Bad Events

We define the following $\mathsf{poly}(k)$ "bad events" ($C_1 > 1$ and C_2 are constants to be determined):

(A) For each non-terminal v in G, let its distance to the its nearest terminal be D_v, and let t be a terminal which is *far* from v; precisely, $\mathsf{dist}_G(v,t) \geq C_1 \cdot D_v$. Let $\mathcal{E}^{\mathsf{far}}_{v,t}$ denote the "bad event" that v is assigned to t.

(B) For each non-terminal v in G, let $\mathcal{E}^{\mathsf{early}}_{v}$ denote the "bad event" that v is assigned before the end of Round-$(C_2 D_v \delta / \log k)$.

We will define another family of "bad events", but before doing so, we need the following setups, which are modified from KKN.[2]

For each pair of terminals t_i, t_j, let the vertices on $\mathsf{SP}(t_i,t_j)$, when listed in the direction from t_i to t_j, be $v_0 = t_i, v_1, v_2, \cdots, v_{L-1}, v_L = t_j$. The vertices $v_1, \cdots, v_{L-1}$ are partitioned into sets, while each set contains consecutive vertices of $\mathsf{SP}(t_i,t_j)$. For each such set $Q = \{v_a, v_{a+1}, \cdots, v_b\}$, we say its *internal length* is the distance between v_a and v_b, while its *external length* is the distance between v_{a-1} and v_{b+1}. (It is possible to have $v_a = v_b$; in this case, the internal length is 0.)

We require each Q in the partition satisfies the following requirement: Q contains a non-terminal v such that the internal length of Q is <u>at most</u> $C_2 D_v \delta/(5 \log k)$ but the external length of Q is <u>at least</u> $C_2 D_v \delta/(5 \log k)$. Such partition can be obtained by a simple greedy sweeping along the directed path $\mathsf{SP}(t_i,t_j)$, where we always take $v \equiv v(Q)$ to be the first vertex of each set Q. The algorithm that constructs G_* guarantees that there are at most $\mathcal{O}(k^2)$ vertices on $\mathsf{SP}(t_i,t_j)$, and thus the cardinality of the partition is bounded by the same number. We denote the partition by $\mathcal{Q}(t_i,t_j)$.

For each $Q = \{v_a, v_{a+1}, \cdots, v_b\} \in \mathcal{Q}(t_i,t_j)$, before the execution of the algorithm, all vertices in Q are *active*. We will then iteratively determine how these vertices are turned to *inactive*. On the course of the algorithm, we say that a terminal t *reaches* Q at a moment when some <u>active vertex</u> in Q is assigned to V_t; when this happens, we let

$$q_{\min} := \min \ \{q \mid v_q \text{ is active just before } t \text{ reaches } Q, \text{ and it is assigned to } t \text{ after } t \text{ reaches } Q\};$$

we define $q_{\max}$ analogously. Then <u>all</u> vertices $\{ v_q \mid q \in [q_{\min}, q_{\max}] \}$ are turned to inactive. Also, we let $\mathsf{Term\text{-}Detour}(q_{\min}, q_{\max}) := \mathsf{SP}(v_{q_{\min}}, t) + \mathsf{SP}(t, v_{q_{\max}}) + (v_{q_{\max}}, v_{q_{\max}+1})$ to be the *terminal-detour* created by this t's reaching on Q. To avoid confusion, we make the following remark: when a vertex is assigned, it must be inactive, but it is possible for a vertex to be inactive but remains unassigned.

Now we are ready to define the last family of "bad events": (C_3 is a constant to be determined)

(C) Let $\mathcal{E}^{\mathsf{many}}_{Q}$ denote the "bad event" that at least $C_3 \log k$ distinct terminals reach Q.

[2] Some of our terminologies are same as those used by KKN, but they have slightly different definitions/meanings.

For the convenience of forthcoming discussion, whenever there are two *consecutive* terminal-detours $\mathsf{Term\text{-}Detour}(q_1, q_2)$ and $\mathsf{Term\text{-}Detour}(q_2+1, q_3)$ which involve the same terminal t, then the two terminal-detours are merged into a new one, which is $\mathsf{SP}(v_{q_1}, t) + \mathsf{SP}(t, v_{q_3}) + (v_{q_3}, v_{q_3+1})$. By repeating this merging process until no consecutive terminal-detours involve the same terminal, we are sure that for any two consecutive terminal-detours $\mathsf{Term\text{-}Detour}(q_1, q_2)$ and $\mathsf{Term\text{-}Detour}(q_2 + 1, q_3)$, the two vertices v_{q_2}, v_{q_2+1} are assigned to two different terminals, and the two terminals are adjacent in the graph G' (due to the edge $(v_{q_2}, v_{q_2+1}) \in E$).

3.3 Key Lemmas and Proof of Theorem 1

By suitably choosing the constants C_1, C_2, C_3, we can show that the above $\mathsf{poly}(k)$ bad events each occurs with probability at most $\mathcal{O}(1/k^\beta)$ for some sufficiently large β. Then by a simple union bound, we can show that with probability $1 - \mathcal{O}(1/k)$, all these bad events are avoided. As we will show in the proof of Theorem 1, the distortion of any terminal-centered minor that avoids all bad events is at most $1 + \frac{40C_3(C_1+1)}{C_2} \log^2 k \ = \ \mathcal{O}(\log^2 k)$.

Lemma 2. *Set $C_1 = 5400$. Then $\mathbb{P}\left[\ \mathcal{E}^{\mathsf{far}}_{v,t}\ \right] \ \leq \ 3/k^6$.*

Lemma 3. *Set $C_2 = 1/27$. Then $\mathbb{P}\left[\ \mathcal{E}^{\mathsf{early}}_{v}\ \right] \ \leq \ 2/k^5$.*

Lemma 4. *Set $C_3 = 30$. Then $\mathbb{P}\left[\ \mathcal{E}^{\mathsf{many}}_{P} \mid \forall v,\ \mathcal{E}^{\mathsf{early}}_{v} \text{ does not occur}\ \right] \ \leq \ 1/k^5$.*

The proofs of the above three lemmas will be provided in Section 4. We follow KKN's method closely to prove Lemma 4. To prove Lamma 2 and Lemma 3, we use the following two lemmas about the tail bounds of sum of independent ERV; their proofs will be presented in Section 5. We note that sum of independent ERV with general *distinct* means follows *hypoexponential distribution*, for which complicated tail bounds are known, but none of them is simple enough for us to use directly. Thus, we derive simpler tail bounds ourselves. This is done by a proper translation to finding tail bound of sum of independent ERV with equal means, which is known to follow an *Erlang distribution* and admits a simpler tail bound.

Lemma 5. *For any positive integer m, let $E_1, E_2, E_3, \cdots, E_m$ be m independent exponential random variables, each has mean at least $A > 0$ (the means can be distinct). Then for any $\kappa \leq 1/4$,*

$$\mathbb{P}\left[\ \sum_{i=1}^{m} E_i \ \leq \ \kappa \cdot Am\ \right] \ \leq \ \frac{4}{3\sqrt{2\pi m}} \cdot (3\kappa)^m \ \leq \ (3\kappa)^m.$$

Lemma 6. *Let $R_1, R_2, R_3, \cdots$ be a sequence of independent exponential random variables with means $A, A/r, A/r^2, \cdots$, where $A > 0$, and $r = 1 + \delta/\log k$ with $\delta \leq 1/2$ and $k \geq 3$. For any sufficiently large k, and for any constant $M \geq 18$,*

$$\mathbb{P}\left[\ \sum_{i=1}^{\infty} R_i \ \geq \ M \cdot \frac{A \log k}{\delta}\ \right] \ \leq \ \frac{2}{k^{M/(12\delta)+3}}.$$

Note that in Lemma 6, the expected value of $\sum_{i=1}^{\infty} R_i$ is $\frac{Ar}{r-1} \ \leq \ \frac{2A\log k}{\delta}$. In other words, the lemma states that the probability for the sum exceeding M times of its expected value is at most $k^{-\Theta_\delta(M)}$.

The intuition behind proof of Lemma 2 is: when t is *far away* from v but there are some other terminals *near* v, it looks highly likely that v is assigned (to the one of the near terminals) well before t grows its ball radius large enough to come close to v. We will verify this natural intuition rigorously. Formally, we design a random variable to track how v is being approached by the sets in the terminal-centered partition; the random variable has initial value zero, and its value is at least D_v when v is assigned. By showing that the random variable dominates some sum of independent ERV, we prove that with high probability, v is assigned early enough (using Lemma 5) while R_t is still below $C_1 \cdot D_v$ (using Lemma 6).

Lemma 3 is almost a direct corollary of Lemma 6.

3.4 Proof of Theorem 1

We finish this technical overview by proving Theorem 1 using Lemmas 2—4.

Let the terminals be $t_1, \cdots, t_k$ and let $V_1, \cdots, V_k$ be a terminal-centered partition which avoids all the above bad events — such parition exists for any sufficiently large k, since by Lemmas 2—4, the probability that any of the bad events occur is at most

$$\mathcal{O}(k^5) \cdot \frac{3}{k^6} \ + \ \mathcal{O}(k^4) \cdot \frac{2}{k^5} \ + \ \binom{k}{2} \cdot \mathcal{O}(k^2) \cdot \frac{1}{k^5} \ = \ \mathcal{O}\left(\frac{1}{k}\right) \ \ll \ 1.$$

For each pair of terminals t_i, t_j, the length of $\mathsf{SP}(t_i, t_j)$ is bounded below by half of the external lengths of all sets in $\mathcal{Q}(t_i, t_j)$ (since some of the edges on $\mathsf{SP}(t_i, t_j)$ are double counted, we need the factor of half), i.e.,

$$\mathsf{dist}_G(t_i, t_j) \ \geq \ \frac{1}{2} \sum_{Q \in \mathcal{Q}(t_i, t_j)} \frac{C_2 D_{v(Q)} \delta}{5 \log k}.$$

Recall that the vertices on $\mathsf{SP}(t_i, t_j)$, when listed in the direction from t_i to t_j, are $v_0 = t_i, v_1, v_2, \cdots, v_{L-1}, v_L = t_j$. For each $Q = \{v_a, v_{a+1}, \cdots, v_b\} \in \mathcal{Q}(t_i, t_j)$, note that conditioning on the avoidance of $\mathcal{E}_Q^{\mathsf{many}}$, there is a directed path from v_a to v_{b+1} in G which is composed of at most $C_3 \log k$ terminal-detours; we denote this directed path by P_Q.

Concatenating (t_i, v_1) with all P_Q, where Q runs over $\mathcal{Q}(t_i, t_j)$, forms a directed path from t_i to t_j in G, denoted by $\mathsf{Detour}(t_i, t_j)$.[3] Due to the terminal-detour-merging process described at the end of Section 3.2, every two consecutive terminal-detours involve two different terminals, and the two terminals are adjacent in G'. Thus $\mathsf{Detour}(t_i, t_j)$, when contracted according to the underlying terminal-center partition, forms a path from t_i to t_j in G'. This path in G' has length bounded above by the length of $\mathsf{Detour}(t_i, t_j)$ in G, due to the definition of w' in $G' = (T, E', w')$. Conditioned on the avoidances of all the bad events, the length of $\mathsf{Detour}(t_i, t_j)$ is upper bounded by

$$\mathsf{dist}_G(t_i, t_j) \ + \ \sum_{Q \in \mathcal{Q}(t_i, t_j)} \left[2(C_1 + 1) \cdot D_{v(Q)}\right] \cdot (C_3 \log k),$$

and hence the distortion is bounded above by

$$1 + \frac{\sum_{Q \in \mathcal{Q}(t_i, t_j)} \left[2(C_1 + 1) \cdot D_{v(Q)}\right] \cdot (C_3 \log k)}{\frac{1}{2} \ \sum_{Q \in \mathcal{Q}(t_i, t_j)} \frac{C_2 D_{v(Q)} \delta}{5 \log k}} \ \leq \ 1 + \frac{40 C_3 (C_1 + 1)}{C_2} \log^2 k,$$

which is at most $(2 \times 10^8) \cdot \log^2 k$.

4 Analysis

4.1 Proof of Lemma 2

Setup. Let t_* denote the nearest terminal to v in graph G, and let P denote the *directed* shortest path from t_* to v. By definition, the length of P is D_v. We use the variable u_* to track the vertex in P which is assigned and furthest away from t_*. At the beginning, the only assigned vertex in P is t_*, thus $u_* = t_*$; when $u_* = v$, v is assigned.

Let t_q be the variable tracking the terminal which u_* is assigned to. Let d_q denote the distance between t_q and u_* in $G[V_q]$, then consider the following random variable:

$$d_* \ := \ \mathsf{dist}_G(t_*, u_*) \ + \ R_q \ - \ d_q.$$

[3] Note that this directed path may not be *simple*, i.e., it may traverse a terminal or an edge for multiple times. This does not matter for us, since the purpose for introducing this path is to give an over-estimate of $\mathsf{dist}_{G'}(t_i, t_j)$.

Note that d_* never decreases, and $d_* \geq D_v$ if and only if v is assigned.

Step 1. In this step, we show that after Round-T, where $T = 200D_v\delta/\log k$, $d_* \geq D_v$ occurs with probability at least $1 - 1/k^6$.

There are two scenarios in which d_* increases: (a) when R_q increases, while u_* might remain unchanged or not; or (b) when some vertex on P which is further away from t_* than u_* is assigned to some terminal $t_{q'}$, where $q' \neq q$.

In scenario (a), d_* increases by exactly the amount of the ERV that raises R_q.

In scenario (b), since u_* is changed, to avoid confusion, we use u_*^{old} to denote the value of u_* before the change. Let u' denote the vertex on the directed path $\mathsf{SP}(u_*^{\text{old}}, v)$ which is newly assigned to $t_{q'}$ and nearest to $t_{q'}$ in $G[V_{q'}]$. Due to the memoryless property of ERV, the increment of d_* is the distance between u_*^{old} and u', plus a random variable that follows an exponential distribution with the same mean as the ERV that raises $R_{q'}$.

On the other hand, observe that in each round, at least one of the two scenarios occurs (and for (b), it might occur multiple times in one round). Combined with the last paragraph, after Round-T, the value of d_* is lower bounded by the sum of independent ERV of the form $\sum_{i=1}^{m} \mathsf{Exp}(T/r^{i-1})$, where $m = \lceil (\log k)/\delta \rceil$. Note that the mean of $T/r^{m-1} \geq T/3 = 200D_v\delta/(3\log k)$. Thus, by Lemma 5,

$$\mathbb{P}\left[\, d_* < D_v \,\right] \;\leq\; \mathbb{P}\left[\sum_{i=1}^{m} \mathsf{Exp}(T/r^{i-1}) \;<\; \frac{3}{200} \cdot \frac{200D_v\delta}{3\log k} \cdot m \right] \;\leq\; \left(\frac{9}{200}\right)^{\log k/\delta} \;\leq\; \frac{1}{k^6}.$$

Step 2. Conditioned on the highly probable event analyzed in Step 1, for the event $\mathcal{E}_{v,t}^{\mathsf{far}}$ to occur, a necessary (but not sufficient) condition is: the radius R_t grows beyond $C_1 D_v = 5400D_v$ by the end of Round-T. By Lemma 6, for sufficiently large k, this occurs with probability at most

$$\mathbb{P}\left[\sum_{i=0}^{\infty} \mathsf{Exp}\left(\frac{3T}{2}/r^i\right) \;\geq\; 18 \cdot \frac{300D_v\delta}{\log k} \cdot \frac{\log k}{\delta} \right] \;\leq\; \frac{2}{k^6}.$$

Thus, we can conclude that $\mathbb{P}\left[\, \mathcal{E}_{v,t}^{\mathsf{far}} \,\right] \;\leq\; 1/k^6 + 2/k^6 \;=\; 3/k^6$.

4.2 Proof of Lemma 3

Lemma 3 can be proved almost directly from Lemma 6. For the event $\mathcal{E}_v^{\mathsf{early}}$ to occur, a necessary (but not sufficient) condition is: there exists a terminal t with radius R_t grows beyond D_v by the end of Round-$(C_2 D_v\delta/\log k)$. By Lemma 6, this occurs with probability at most

$$\begin{aligned}
&\sum_{t\in T} \mathbb{P}\left[\sum_{i=0}^{\infty} \mathsf{Exp}\left(\frac{3C_2D_v\delta}{2\log k}/r^i\right) \;\geq\; D_v \right] \\
&\qquad =\; k \;\cdot\; \mathbb{P}\left[\sum_{i=0}^{\infty} \mathsf{Exp}\left(\frac{3C_2D_v\delta}{2\log k}/r^i\right) \;\geq\; \frac{2}{3C_2} \cdot \frac{3C_2D_v\delta}{2\log k} \cdot \frac{\log k}{\delta} \right] \;\leq\; k \;\cdot\; \frac{2}{k^6} \;=\; \frac{2}{k^5}.
\end{aligned}$$

4.3 Proof of Lemma 4

The idea behind the proof of Lemma 4 follows closely from [KKN15, Sections 3.1.1 and 3.1.2], but there are some small differences in the details. For completeness, we provide a self-contained proof.

Let $Q = \{v_a, v_{a+1}, \cdots, v_b\}$. On the course of the algorithm, a *maximal active subset* of Q is a subset $Q' = \{v_q, v_{q+1}, \cdots, v_r\} \subset Q$ such that (a) all vertices in Q' are active; (b) v_{q-1} is inactive or $q = a$; and (c) v_{r+1} is inactive or $r = b$. We let Z_Q count the number of maximal active subsets of Q on the course of the algorithm. At the beginning, the only maximal active subset of Q is Q itself, so

$Z_Q = 1$. Whenever $Z_Q = 0$, we obtain a collection of terminal-detours which are later used to form P_Q, which we have defined in the proof of Theorem 1.

For each $\ell \in \mathbb{N}$, let t^ℓ denote the ℓ-th terminal that reaches Q, then let Y_ℓ denote the indicator random variable that upon the reach of t^ℓ on Q, Z_Q strictly decreases. We note that when $Y_\ell = 0$, by inspecting the process of how vertices are turned inactive (see Section 3.2), it is not hard to see that Z_Q increases by at most 1, and furthermore, it must be the case that *some but not all* vertices in one maximal active subset get assigned to t^ℓ.[4]

Note that for any vertex $v \in Q$, $D_v \ \geq \ D_{v(Q)} \cdot (1 - C_2\delta/5\log k) \ \geq \ 99D_{v(Q)}/100$. Conditioned on the avoidances of all $\mathcal{E}_v^{\mathsf{early}}$, a vertex in Q can be assigned only after Round-$(99C_2D_{v(Q)}\delta/(100\log k))$. Also, note that the internal length of each maximal active subset is at most the internal length of Q, which is at most $C_2D_{v(Q)}\delta/(5\log k)$. Hence, by the memoryless property of ERV, we have

$$\mathbb{P}\left[\ Y_\ell = 0\ \right] \ \leq \ \mathbb{P}\left[\ \mathsf{Exp}\left(\frac{99C_2D_{u(Q)}\delta}{100\log k}\right) \ \leq \ \frac{C_2D_{u(Q)}\delta}{5\log k}\ \right] \ \leq \ \frac{20}{99}.$$

Next, for an integer N to be determined, let $\overline{Y}_N := \sum_{\ell=1}^{N} Y_\ell$. Note that when $\overline{Y}_N > N/2$, then $Z_Q = 0$. Thus, $\mathbb{P}\left[\ Z_Q \geq 1 \text{ after } Q \text{ being reached for } N \text{ times }\right] \ \leq \ \mathbb{P}\left[\ \overline{Y}_N \leq N/2\ \right]$, which is bounded above by $\mathbb{P}\left[\ \mathsf{Binom}(N, 79/99) \ \leq \ N/2\ \right]$.[5] By the Hoeffding's inequality,

$$\mathbb{P}\left[\ \mathsf{Binom}\left(N, \frac{79}{99}\right) \ \leq \ \frac{N}{2}\ \right] \ \leq \ e^{-2(79/99-1/2)^2 N}.$$

By choosing $N = 30\log k$, the RHS of the above inequality is at most $1/k^5$.

5 Tail Bounds of Sum of Exponential Random Variables

In this section, we provide proofs of Lemma 5 and Lemma 6.

Proof of Lemma 5: First, since ERV enjoys the property of closeness under scaling, it suffices to prove the lemma for $A = 1$.

It is well-known that sum of independent ERV with *equal* means follows an Erlang distribution. Thus,

$$\mathbb{P}\left[\ \sum_{i=1}^{m} E_i \ \leq \ \kappa m\ \right] \ \leq \ \mathbb{P}\left[\ \sum_{i=1}^{m} \mathsf{Exp}(1) \ \leq \ \kappa m\ \right] \ = \ \mathbb{P}\left[\ \mathsf{Erlang}(m,1) \ \leq \ \kappa m\ \right] \ = \ \frac{\gamma(m,\kappa m)}{(m-1)!},$$

where $\gamma(\cdot,\cdot)$ is the *lower incomplete Gamma function*, which admits a power series expansion given below [Par]. Together with the assumption $\kappa \leq 1/4$, it yields

$$\begin{aligned}\frac{\gamma(m,\kappa m)}{(m-1)!} \ &= \ \frac{e^{-\kappa m}(\kappa m)^m}{(m-1)!} \cdot \sum_{\ell=0}^{\infty} \frac{(\kappa m)^\ell}{m(m+1)\cdots(m+\ell)} \\ &\leq \ \frac{e^{-\kappa m}(\kappa m)^m}{m!} \cdot \left[1 + \frac{1}{4} + \frac{1}{16} + \cdots\right] \ = \ \frac{4e^{-\kappa m}(\kappa m)^m}{3 \cdot m!}.\end{aligned}$$

[4] An elaborative explaination: (1) If $v_{q_{\min}}$ and $v_{q_{\max}}$ belong two different maximal active subsets (denoted as S_1, S_2), then all maximal active subsets strictly between S_1, S_2 are turned to inactive, while S_1, S_2 might shrink to a smaller maximal active subset or be turned to fully inactive. When either S_1, S_2 is turned to fully inactive, Z_Q must drop by at least 1. Thus, whenever Z_Q does not strictly decrease, some vertices in either S_1, S_2 remain unassigned. (2) If $v_{q_{\min}}$ and $v_{q_{\max}}$ belong to the same maximal active subset S_1, then it's clear that Z_Q can increase by at most 1, and when Z_Q strictly increases, some vertices in S_1 remain unassigned.

[5] $\mathsf{Binom}(N,p)$ is the Binomial distribution followed by $\sum_{\ell=1}^{N} X_\ell$, where each X_ℓ is an independent random variable that takes value 1 with probability p, and it takes value 0 otherwise.

By the Stirling's approximation, $m! \geq \sqrt{2\pi m} \cdot (m/e)^m$, thus we can complete the proof:

$$\frac{4e^{-\kappa m}(\kappa m)^m}{3 \cdot m!} \;\leq\; \frac{4}{3\sqrt{2\pi m}} \cdot [e^{1-\kappa}\kappa]^m \;\leq\; \frac{4}{3\sqrt{2\pi m}} \cdot (3\kappa)^m.$$

□

To prove Lemma 6, we need the following intermediate lemma.

Lemma 7. *Let $E_1, E_2, E_3, \cdots, E_m$ be m independent exponential random variables, each has mean at most $A > 0$. Then for any sufficiently large m, and for any $C \geq 6$,*

$$\mathbb{P}\left[\; \sum_{i=1}^{m} E_i \;\geq\; C \cdot Am \;\right] \;\leq\; \frac{2}{(C-1)\cdot\sqrt{2\pi m}} \cdot \frac{1}{e^{Cm/2}} \;\leq\; \frac{1}{e^{Cm/2}}.$$

Proof: Again, it suffices to prove the lemma for $A = 1$. Note that

$$\mathbb{P}\left[\; \sum_{i=1}^{m} E_i \;\geq\; Cm \;\right] \;\leq\; \mathbb{P}\left[\; \sum_{i=1}^{m} \mathsf{Exp}(1) \;\geq\; Cm \;\right] \;=\; \mathbb{P}\left[\; \mathsf{Erlang}(m,1) \;\geq\; Cm \;\right] \;=\; \frac{\Gamma(m, Cm)}{(m-1)!},$$

where $\Gamma(\cdot,\cdot)$ is the *upper incomplete Gamma function*. Tricomi [Tri50] (see also [Gau98]) showed that for any constant $C \geq 6$,

$$\Gamma(m, Cm) \;=\; \frac{e^{-Cm}\;(Cm)^m}{Cm - m + 1} \cdot (1 + o_m(1)).$$

By the Stirling's approximation, $m! \geq \sqrt{2\pi m} \cdot (m/e)^m$. Thus, for any sufficiently large m,

$$\frac{\Gamma(m, Cm)}{(m-1)!} \;\leq\; \frac{2e^{-Cm}(Cm)^m}{(C-1)\cdot m!} \;\leq\; \frac{2}{(C-1)\cdot\sqrt{2\pi m}} \cdot [e^{1-C}C]^m.$$

When $C \geq 6$, $e^{1-C}C \leq e^{-C/2}$, and we are done. □

Proof of Lemma 6: Again, it suffices to prove the lemma for $A = 1$. Let $m := \lfloor (\log k)/\delta \rfloor \;\geq\; \frac{\log k}{2\delta}$. For $j \in \mathbb{N}$, let

$$\mathcal{E}_j \text{ denote the following event: } \sum_{r=(j-1)m+1}^{jm} R_r \;\geq\; \frac{1}{2^{j-1}} \cdot \left(\frac{1}{2} + \frac{j}{2}\right) \cdot \frac{M \log k}{3\delta}.$$

Due to the equality $\sum_{j=1}^{\infty}(1+j)/2^j \;=\; 3$, if $\sum_{i=1}^{\infty} R_i \;\geq\; \frac{M \log k}{\delta}$, then at least one of $\mathcal{E}_1, \mathcal{E}_2, \cdots$ must hold. Hence,

$$\mathbb{P}\left[\; \sum_{i=1}^{\infty} R_i \;\geq\; \frac{M \log k}{\delta} \;\right] \;\leq\; \sum_{j=1}^{\infty} \mathbb{P}\left[\; \mathcal{E}_j \;\right].$$

Next, observe that $\mathbb{E}\left[R_{(j-1)m+1}\right] \;\leq\; (1 + 1/m)^{-(j-1)m} \;\leq\; 1/2^{j-1}$. Hence, by Lemma 7,

$$\mathbb{P}\left[\; \mathcal{E}_j \;\right] \;\leq\; 1 \Big/ \exp\left[\left(\frac{1}{2} + \frac{j}{2}\right) \cdot \frac{M \log k}{3\delta} \cdot \frac{1}{2}\right] \;=\; k^{-(1+j)M/(12\delta)} \;\leq\; \frac{1}{k^{M/(12\delta)}} \cdot \frac{1}{k^{3j}},$$

and thus

$$\mathbb{P}\left[\; \sum_{i=1}^{\infty} R_i \;\geq\; \frac{M \log k}{\delta} \;\right] \;\leq\; \frac{1}{k^{M/(12\delta)+3}} \left(1 + \frac{1}{k^3} + \frac{1}{k^6} + \cdots\right) \;\leq\; \frac{2}{k^{M/(12\delta)+3}}.$$

□

6 Discussion

If we are allowed to use only one brief sentence to summarize how the $\mathcal{O}(\log^2 k)$ distortion is arrived, it is: the vertices in a short segment of length $\ell = \Theta(D_v/\log k)$ might be assigned to at most $\Theta(\log k)$ distinct terminals, each is of distance of at most $\Theta(\ell \log k) = \Theta(D_v)$ away from the segment.

In our (failed) attempts to shave further log, one idea is to grow the ball radius more aggressively by setting $r = 1 + \Theta(1)$ instead of $r = 1 + \Theta(\log^{-1} k)$. To motivate this idea, see in the proof of Lemma 2 and by Lemma 3, a vertex v is assigned in Round-$[\Theta(D_v/\log k)]$. This forced us to set ℓ to be of the same magnitude for establishing Lemma 4. By setting $r = 1 + \Theta(1)$, if the sum of ERV $\sum_{i=1}^{\infty} \mathsf{Exp}(\Theta(D_v)/r^i)$ were well concentrated around its mean, we could hope that v is assigned in Round $\Theta(D_v)$ and then we could set $\ell = \Theta(D_v)$. Unfortunately, such sum of ERV is actually *not* sufficiently well concentrated.

Another idea is to replace the exponential random variables, which have unbounded support, with some other random variables, say those with bounded supports. Then Chernoff-like tail bounds might be plausible. However, memoryless property is lost. While I have my personal intuition that this approach might provide better results, the loss of memoryless property makes devising a rigorous proof seem technically challenging.

This draws a parallel with the early development of *queueing theory*. Its fundamental results mostly apply to models in which the involved random variables are ERV — the reason is not because ERV is the most appropriate probability distribution that captures the reality, but because the memoryless property of ERV makes analyses more admissible. After more than a century of development, rich analytical insights and tools have been developed to handle more general queueing models; it might be interesting to investigate whether such insights and tools can help us to "grow balls for the better".

References

[AGK14] Alexandr Andoni, Anupam Gupta, and Robert Krauthgamer. Towards $(1 + \epsilon)$-approximate flow sparsifiers. In *SODA 2014*, pages 279–293, 2014.

[BG08] Amitabh Basu and Anupam Gupta. Steiner point removal in graph metrics. 2008. http://www.ams.jhu.edu/∼abasu9/papers/SPR.pdf.

[CGH16] Yun Kuen Cheung, Gramoz Goranci, and Monika Henzinger. Graph minors for preserving terminal distances approximately - lower and upper bounds. In *ICALP 2016*, pages 131:1–131:14, 2016.

[Chu12] Julia Chuzhoy. On vertex sparsifiers with steiner nodes. In *STOC 2012*, pages 673–688, 2012.

[CLLM10] Moses Charikar, Tom Leighton, Shi Li, and Ankur Moitra. Vertex sparsifiers and abstract rounding algorithms. In *FOCS 2010*, pages 265–274, 2010.

[CXKR06] Hubert T.-H. Chan, Donglin Xia, Goran Konjevod, and Andréa W. Richa. A tight lower bound for the steiner point removal problem on trees. In *APPROX 2006*, pages 70–81, 2006.

[EFN15a] Michael Elkin, Arnold Filtser, and Ofer Neiman. Prioritized metric structures and embedding. In *STOC 2015*, pages 489–498, 2015.

[EFN15b] Michael Elkin, Arnold Filtser, and Ofer Neiman. Terminal embeddings. In *APPROX 2015*, pages 242–264, 2015.

[EGK+14] Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Räcke, Inbal Talgam-Cohen, and Kunal Talwar. Vertex sparsifiers: New results from old techniques. *SIAM J. Computing*, 43(4):1239–1262, 2014.

[Gau98] Walter Gautschi. The incomplete gamma functions since Tricomi. In *Tricomi's Ideas and Contemporary Applied Mathematics, Atti dei Convegni Lincei, n. 147, Accademia Nazionale dei Lincei*, pages 203–237, 1998.

[Gup01] Anupam Gupta. Steiner points in tree metrics don't (really) help. In *SODA 2001*, pages 220–227, 2001.

[KKN15] Lior Kamma, Robert Krauthgamer, and Huy L. Nguyen. Cutting corners cheaply, or how to remove steiner points. *SIAM J. Computing*, 44(4):975–995, 2015.

[KNZ14] Robert Krauthgamer, Huy L. Nguyen, and Tamar Zondiner. Preserving terminal distances using minors. *SIAM J. Discrete Math.*, 28(1):127–141, 2014.

[KR13] Robert Krauthgamer and Inbal Rika. Mimicking networks and succinct representations of terminal cuts. In *SODA 2013*, pages 1789–1799, 2013.

[LM10] Tom Leighton and Ankur Moitra. Some results on greedy embeddings in metric spaces. *Discrete & Computational Geometry*, 44(3):686–705, 2010.

[MM16] Konstantin Makarychev and Yury Makarychev. Metric extension operators, vertex sparsifiers and lipschitz extendability. *Israel Journal of Mathematics*, 212(2):913–959, 2016.

[Moi09] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *FOCS 2009*, pages 3–12, 2009.

[Par] Richard B. Paris. Incomplete gamma and related functions. http://dlmf.nist.gov/8.

[RTZ05] Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *ICALP 2005*, pages 261–272, 2005.

[Tri50] F. Tricomi. Asymptotische eigenschaften der unvollständigen gammafunktion. *Mathematische Zeitschrift*, 53:136–148, 1950.